International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

Q.3313
(02/2012)

SERIES Q: SWITCHING AND SIGNALLING

Signalling requirements and protocols for the NGN –
Resource control protocols

Signalling protocols and procedures relating to flow state aware QoS control in a bounded subnetwork of a next generation network

Recommendation ITU-T Q.3313

ITU-T

International Telecommunication Union

ITU-T Q-SERIES RECOMMENDATIONS

SWITCHING AND SIGNALLING

SIGNALLING IN THE INTERNATIONAL MANUAL SERVICE	Q.1–Q.3
INTERNATIONAL AUTOMATIC AND SEMI-AUTOMATIC WORKING	Q.4–Q.59
FUNCTIONS AND INFORMATION FLOWS FOR SERVICES IN THE ISDN	Q.60–Q.99
CLAUSES APPLICABLE TO ITU-T STANDARD SYSTEMS	Q.100–Q.119
SPECIFICATIONS OF SIGNALLING SYSTEMS No. 4, 5, 6, R1 AND R2	Q.120–Q.499
DIGITAL EXCHANGES	Q.500–Q.599
INTERWORKING OF SIGNALLING SYSTEMS	Q.600–Q.699
SPECIFICATIONS OF SIGNALLING SYSTEM No. 7	Q.700–Q.799
Q3 INTERFACE	Q.800–Q.849
DIGITAL SUBSCRIBER SIGNALLING SYSTEM No. 1	Q.850–Q.999
PUBLIC LAND MOBILE NETWORK	Q.1000–Q.1099
INTERWORKING WITH SATELLITE MOBILE SYSTEMS	Q.1100–Q.1199
INTELLIGENT NETWORK	Q.1200–Q.1699
SIGNALLING REQUIREMENTS AND PROTOCOLS FOR IMT-2000	Q.1700–Q.1799
SPECIFICATIONS OF SIGNALLING RELATED TO BEARER INDEPENDENT CALL CONTROL (BICC)	Q.1900–Q.1999
BROADBAND ISDN	Q.2000–Q.2999
SIGNALLING REQUIREMENTS AND PROTOCOLS FOR THE NGN	Q.3000–Q.3999
General	Q.3000–Q.3029
Network signalling and control functional architecture	Q.3030–Q.3099
Network data organization within the NGN	Q.3100–Q.3129
Bearer control signalling	Q.3130–Q.3179
Signalling and control requirements and protocols to support attachment in NGN environments	Q.3200–Q.3249
Resource control protocols	**Q.3300–Q.3369**
Service and session control protocols	Q.3400–Q.3499
Service and session control protocols – supplementary services	Q.3600–Q.3649
NGN applications	Q.3700–Q.3849
Testing for NGN networks	Q.3900–Q.3999

For further details, please refer to the list of ITU-T Recommendations.

Recommendation ITU-T Q.3313

Signalling protocols and procedures relating to flow state aware QoS control in a bounded subnetwork of a next generation network

Summary

Recommendation ITU-T Q.3313 specifies the signalling format and procedures for the flow state aware (FSA) transfer capability in a bounded subnetwork of a next generation network (NGN). The FSA transfer capability provides QoS controls that operate on a per-flow basis, allowing flows within a bounded subnetwork of an NGN to receive different treatments depending on signalled parameters. These parameters are requested using in-band signalling. The parameters contained in these signals are included in the "flow state" maintained on each flow (or each aggregate flow), at each FSA node.

Service options that may be selected include the requested support of the highest available end-to-end rate for data transfer, assuming some source-to-receiver forwarding paths are entirely within the bounded subnetwork of an NGN.

The focus of this Recommendation is on broadband (including mobile) service access scenarios typically involving restricted bandwidth shared by many flows. In such circumstances, customer satisfaction, when there is temporary congestion, may be best handled by applying flow preferences and QoS differently for each customer and not simply on the type of media. This leads to the notion of customised QoS supported partly by signalling and partly through web-based tools.

These concepts may also be applied to flow aggregates, with associated signalling support acting at the aggregate level between aggregation end-points, as defined in Recommendation ITU-T Y.2121. The customization of aggregates in this Recommendation is limited to the notions of preference priority and FSA transfer capability.

History

Edition	Recommendation	Approval	Study Group
1.0	ITU-T Q.3313	2012-02-22	11

Keywords

Available rate, broadband service access, flow state, flow state aware (FSA), in-band signalling, NGN, QoS.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at http://www.itu.int/ITU-T/ipr/.

Table of Contents

		Page
1	Scope	1
2	References	1
3	Definitions	1
	3.1 Terms defined elsewhere	1
	3.2 Terms defined in this Recommendation	2
4	Abbreviations and acronyms	3
5	Conventions	4
6	High-level description	4
7	Protocol description	5
	7.1 QoS parameters related to flow-based QoS management	5
	7.2 Message sequences related to FSA transport service set up and clear down.	6
	7.3 Procedures and format of layer 2 FSA Ethertype packets	8
8	Security considerations and requirements	12
	8.1 Authentication	12
	8.2 Authorization	12
	8.3 Data confidentiality	13
	8.4 Data integrity	13
	8.5 Accountability	13
	8.6 Availability and accessibility	13
	8.7 Privacy	13
	8.8 Protection against network attacks, from within or outside	13
9	State diagrams for FSA signalling edge functions at origination and destination	13
	9.1 Origination state machine	13
	9.2 Termination state machine	14
	9.3 Signalling renegotiation	15
10	FSA QoS manager	15
Annex A – Rules for encoding floating point rates		17
	A.1 AR and GR – Encoding floating point	17
Appendix I – Mapping FSA services to ITU-T Y.1221 transfer capabilities		18
	I.1 Maximum rate service (MRS)	18
	I.2 Available rate service (ARS)	18
Appendix II – Example of a bounded subnet isolated by proxies		20

Recommendation ITU-T Q.3313

Signalling protocols and procedures relating to flow state aware QoS control in a bounded subnetwork of a next generation network

1 Scope

This Recommendation provides the signalling format and procedures for the flow state aware (FSA) transfer capability in a bounded subnetwork of a next generation network (NGN). The FSA transfer capability provides QoS controls that operate on a per-flow basis, allowing flows within a bounded subnetwork of an NGN to receive different treatments depending on signalled parameters. These parameters are requested using in-band signalling. The parameters contained in these signals are included in the "flow state" maintained on each flow (or each aggregate flow), at each FSA node.

Out of scope are any FSA edge functions at the network domain boundaries needed for interworking between two FSA networks, where one uses in-band signalling exclusively for all FSA service support and the other uses both out-of band signalling (for resource reservation and clear down) and in-band signalling (to establish the agreed flow state). FSA signalling must not be used in the core network as it impacts on the scalability.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[ITU-T Y.1221] Recommendation ITU-T Y.1221 (2010), *Traffic control and congestion control in IP-based networks*.

[ITU-T Y.2111] Recommendation ITU-T Y.2111 (2011), *Resource and admission control functions in next generation networks*.

[ITU-T Y.2121] Recommendation ITU-T Y.2121 (2008), *Requirements for the support of flow-state-aware transport technology in NGN*.

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 available rate service (ARS) [ITU-T Y.2121]: The name of one of the flow state aware (FSA) transport services. ARS is primarily for applications that can flexibly adapt to the current available capacity and can quickly adjust their sending rate as the available capacity changes.

3.1.2 flow [ITU-T Y.2121]: A unidirectional sequence of packets with the property that, along any given network link, a flow identifier has the same value for every packet.

3.1.3 flow admission control [ITU-T Y.2121]: The determination, for authorised requests, of whether or not to accept a given flow.

3.1.4 flow aggregate [ITU-T Y.2121]: A hierarchical flow construct that is associated with a group of flows. The carried flows may extend beyond the flow aggregate. Except for the end nodes, flow aggregate forwarders in general do not know that they are carrying flows within the flow aggregate. All packets belonging to a given flow aggregate are commonly routed between aggregation end-points.

3.1.5 flow state [ITU-T Y.2121]: A set of values stored per flow identifier at each flow state aware node. This set of values will determine controls applied on a per-flow basis, dealing with forwarding rate, delay, and congestion recovery.

3.1.6 flow state aware node [ITU-T Y.2121]: A network node that is capable of maintaining flow state and applying per-flow QoS controls, based on recognising flow identifier and associated signals.

3.1.7 in-band signalling [ITU-T Y.2121]: A mode of signalling where the signalling messages are within the flow of the data packets, and follow a path that is tied to the data packets. Signalling messages are routed only through nodes that are in the data path.

3.1.8 maximum rate service (MRS) [ITU-T Y.2121]: The name of one of the flow state aware (FSA) transport services. MRS is for applications that want packet loss characteristics to be sufficient for streamed services as soon as possible but are unwilling to wait or be rejected by network admission control if network resource for this target QoS is not available immediately.

3.1.9 out-of-band signalling [ITU-T Y.2121]: A mode of signalling where the signalling messages are not in the same flow of the data packets, and may follow a different path to the data packets and are routed to one or more nodes that are not in the data path.

3.1.10 QoS structure [ITU-T Y.2121]: The block of QoS signalling information in a signalling packet.

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

3.2.1 bounded subnet: In the context of Recommendation ITU-T Q.3313, a flow state aware (FSA) bounded subnet is a network section where all network nodes support that Recommendation.

3.2.2 ethertype: An Ethertype is a 16 bit identifier assigned by the IEEE to designate a layer 3 protocol being carried over Ethernet.

NOTE – The Ethertype for Recommendation ITU-T Q.3313 is 0x22EF.

3.2.3 flow identifier: A vector comprising the values of a number of elements taken from the header fields of the incoming packets which identify the flow. The flow identifier for a flow within a single FSA network is unique.

3.2.4 flow state aware proxy: A function that provides the origin and/or termination of the flow state aware end-to-end signalling path, and participates in requests and responses on behalf of a user application or management action.

NOTE – Flow state aware proxy may be located, for example, in a user end-system or at a network edge node where it serves as the signalling end-point of multiple users and associated applications.

3.2.5 preference priority: A parameter used to determine whether to admit a flow or scale its rate in case of network overload.

NOTE – In network overload state, an MRS flow with the lower preference priority may be rejected while the one with a higher preference priority level will still be admitted, or for ARS flows, preference priority should be used to scale the rate of the flow with respect to other ARS flows.

3.2.6 proxy: In the context of Recommendation ITU-T Q.3313, proxy is the name for a flow state aware signalling edge function, a process which encapsulates incoming traffic and inserts ITU-T Q.3313 signalling in one direction and in the other direction de-encapsulates and deletes ITU-T Q.3313 signalling.

NOTE – This process can be located inside a user's end-system or exist as a separate system in the data path.

3.2.7 RESPONSETIMEOUT: The maximum period that a signal packet of type request, renegotiate, or close has not been responded to.

NOTE – The signalling packet should be repeated a second and third time before the ITU-T Q.3313 state is dropped.

3.2.8 signal packet: Signal packets are used to carry signalling information across the bounded subnet where flow state aware (FSA) support is provided.

3.2.9 STATETIMEOUT: The maximum period between signalling packets before the network drops a flow.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

AR	Available Rate
ARS	Available Rate Service
ATM	Asynchronous Transfer Mode
BT	Burst Tolerance
CD	Change Direction
CMTS	Cable Modem Termination System
DoS	Denial of Service
FSA	Flow State Aware
FSD	Flow Sender Depth
GBRA	Generic Byte Rate Algorithm
GR	Guaranteed Rate
ISP	Internet Service Provider
MRS	Maximum Rate Service
NGN	Next Generation Network
POP	Point of Presence
PP	Preference Priority
Ptype	Packet Type (in FSA header)
QoS	Quality of Service
SBW	Statistical Bandwidth
TP	Type (of FSA transmission mode)

5 Conventions

ID1	First part of packet ID in FSA header
ID2	Second part of packet ID
ID3	Third part of packet ID – all 3 identify the flow
M	Modified marker
QoS Offset	Byte count to end of packet or start of QoS header in FSA header

6 High-level description

Recommendation [ITU-T Y.2121] specifies the requirements for the support of the flow state aware (FSA) transfer capability in a next generation network (NGN). This transfer capability provides quality of service (QoS) controls that operate on a per-flow basis, allowing flows to receive different treatments depending on their signalled request and the load or congestion at the FSA network nodes. The set of values of all such parameters defines the "flow state" maintained on each flow at each FSA node.

QoS requirements (as would be applied to services supported by FSA transfer) go beyond just the delay and loss that can occur in the transport of incoming packets. The requirements include:

– bandwidth/capacity needed by the application, and

– the priority that bandwidth is maintained during congestion and is restored after various failure events.

To achieve the required QoS for FSA transfer, networks must incorporate the following functions:

1) Functions supporting the FSA packet forwarding behaviours that are applied per flow.

2) Flow admission control recognising and processing requests for associated FSA transport services.

3) Functions supporting the signalling for allocating necessary resources for each flow.

Figure 6-1 shows the main functions which are involved in establishing and ceasing FSA transfer and ensuring the correct provisioning of resources to meet QoS objectives.

Two alternative methods of signalling are shown in Figure 6-1. One method uses in-band signalling exclusively, the other method uses out-of-band signalling for resource reservation and clear down coupled with in-band signals to establish an agreed flow state in each FSA node. Definitions of the terms "in-band signalling" and "out-of-band signalling" are included in [ITU-T Y.2121] to describe the meaning of such terms when linked to the concept of a flow rather than pre-established channels designated for signalling or content transport. The options that are supported are the choice of a network operator or service provider, but any FSA network shall be capable of at least passing every type of FSA in-band signal transparently to allow interoperation with networks that exclusively use such signals.

For the case of in-band signalling, the signalling messages are incorporated into the user data packets themselves, allowing the QoS requirements to be set up during the initial network traversal from sender to receiver (and back if needed). Each FSA node in the path examines the in-band signal and agrees to or adjusts parameter values it can support.

It is assumed that signalling messages originate and terminate at flow state aware signalling edge functions as defined in [ITU-T Y.2121].

As specified in [ITU-T Y.2121], there are four general types of FSA transport services of which two will be used. The first type is a maximum rate flow, which allows some oversubscription but very low delay. The second type is an available rate flow, one that can be used to determine the highest rate the network can immediately support, eliminating slow-start problems.

This Recommendation covers protocol aspects involving both near-end and far-end FSA proxies as well as the FSA node. For out-of-band signals, this Recommendation also includes service control and resource reservation functions as described in [ITU-T Y.2111]. Out-of-band signals are shown by the dotted lines in Figure 6-1.

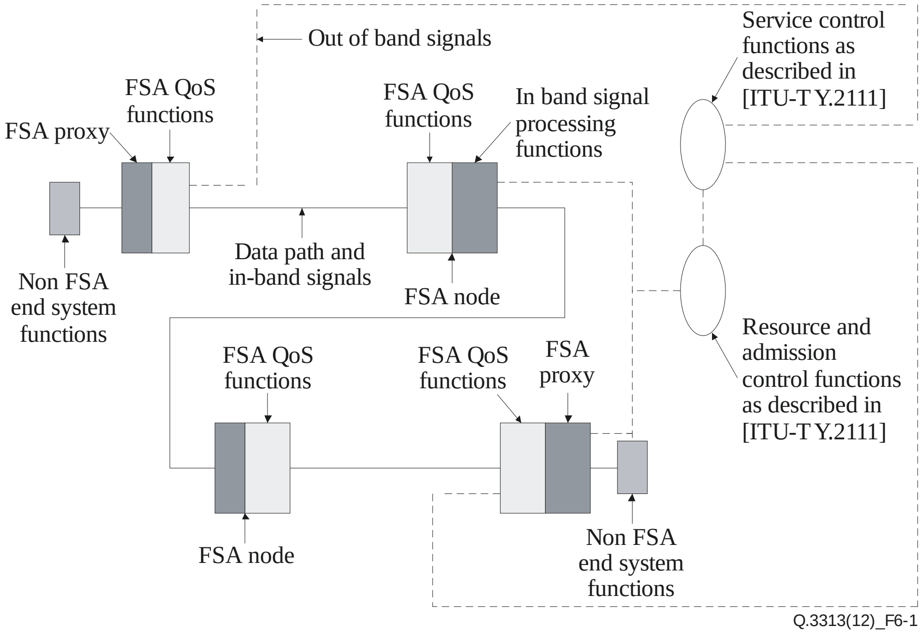

Figure 6-1 – Symmetric access control with FSA-aware end-systems

In Figure 6-1 the communicating FSA proxies are each managed by an FSA node. Furthermore, signalling enables each FSA node to inform flow state information to an upstream FSA node.

Note that there are no routing requirements to consider in any of the cases under consideration in this Recommendation. The upstream forwarding path towards any FSA node may be realised using any appropriate technology that supports the desired QoS.

The FSA node is limited in its functions for supporting QoS on the path between an upstream and a downstream FSA node. It may manage the aggregate traffic forwarded along this path so that it is limited to a given ceiling value, where this given value is provided by network management.

7 Protocol description

This clause details the main model for access QoS management. Parameters that can be requested or modified relating to the QoS treatment of an individual flow or flow aggregate are formulated first.

7.1 QoS parameters related to flow-based QoS management

7.1.1 Overview of the QoS structure

The critical part of the in-band QoS signalling protocol consists of a QoS structure that represents a set of fields containing values and indications on the requested flow treatment or on network responses to this request.

The uses of the QoS structure are as follows:

Travelling in the first packet, the QoS structure is examined by each FSA node to determine if the QoS request can be supported. If it can be supported, the packet proceeds to the next FSA node without change. If the FSA node cannot provide the rates requested, it reduces the requested rate in the QoS structure to what it can support.

If any FSA node finds it cannot continue to accept the rates it has approved, it may selectively discard packets from that flow. This loss of packets is the indication that the receiver may use to create a response towards the source that includes a QoS structure indicating a reduced new rate if appropriate.

Note that the FSA edge QoS manager will need to adjust the available rate (AR) of long duration flows as load builds up.

7.1.2 Overview of QoS structure parameters

Request/Response QoS structure has 10 parameters as follows:

– **AR**: Available rate – floating point rate for network assigned rates.

– **GR**: Guaranteed rate – floating point rate for requested guaranteed rate.

– **PP**: Preference priority – indicates the override or relative rate priority of the flow.

– **CD**: Change/direction – indicating an action type, such as request or response.

– **TP**: Type of FSA transport, such as available rate, or maximum rate.

– **Second QoS structure attached**: Indicates that two QoS structures are included in the same packet. For example, the first one could relate to a response to a forward request and the other one to a separate request for QoS in the reverse direction.

– **Security structure attached**: Indicates that a security structure follows the QoS structure(s).

– **QoS version**: QoS structure version – Set to 2.

– **M**: Modified marker. Set to 0 by sender on request. Set to 1 by FSA nodes if any field changed during a request or renegotiate. Set to 0 and not changed on Response.

– **FSD**: Flow sender depth – Number of proxies entered but not exited.

7.2 Message sequences related to FSA transport service set up and clear down

7.2.1 FSA maximum rate service (MRS), sender requests set-up

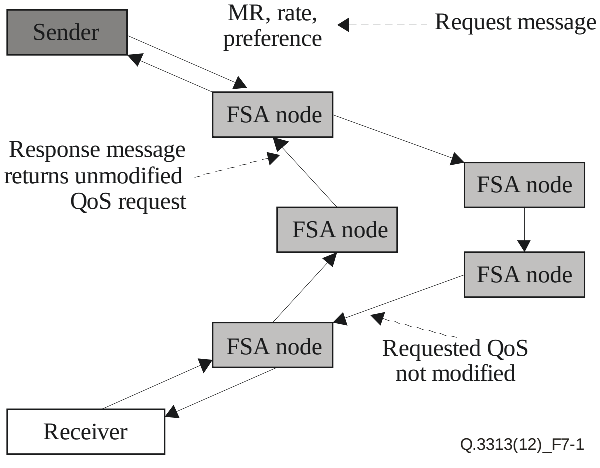

Figure 7-1 – Maximum rate service message sequence

The "Sender" and "Receiver" in Figures 7-1 and 7-2 are the parts of proxies that manage the ITU-T Q.3313 signalling. Figures 7-1 and 7-2 show the signalling process within the ITU-T Q.3313 bounded subnet.

In Figure 7-1, the source (sender) forwards to the receiver a request packet with the requested QoS. This packet is read by intermediate FSA nodes as well as the receiver and includes the following QoS parameters:

– Requested FSA transport service = MRS

– Requested maximum rate = GR

– Requested preference priority = PP

After sending the request, if no response is received within RESPONSETIMEOUT then the request is repeated for up to three times. If all three tries fail, then the FSA signalling should be terminated and the data flow should be sent without any FSA signalling or encapsulation. If a valid response is received, then the data packets may start to be sent at the rate received in the response. If the response has zero as the rate, no data should be sent. The default value of RESPONSETIMEOUT is one second.

7.2.2 MRS clear down

Flow state times out and is terminated if data packets corresponding to that flow are not seen at an FSA node for a period equal to or greater than STATETIMEOUT. The default value for STATETIMEOUT is 2 seconds.

7.2.3 FSA available rate service (ARS), sender requests set-up

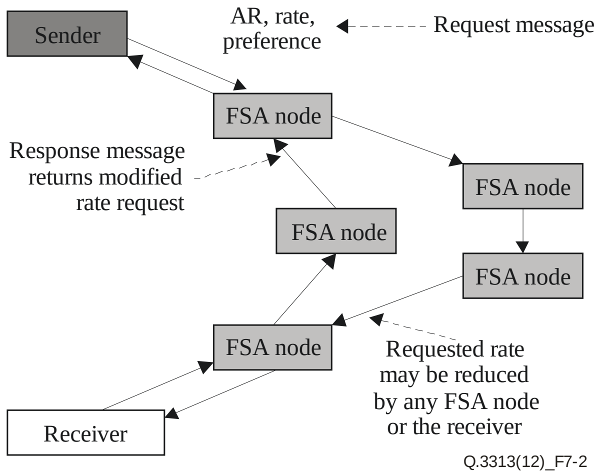

Figure 7-2 – Available rate service message sequence

In Figure 7-2, the sender forwards to the receiver a sequence of request packets. All such packets, after the first request with CD = 1 should be Renegotiate packets with CD = 5. It is a requirement that the sender forwards a new AR renegotiate packet after every 128 data packets or 1 second, whichever is earlier. Each request packet will be read by intermediate FSA nodes as well as by the receiver and must include the following QoS parameters:

– Requested FSA transport service = ARS

– Requested available rate = AR

– Requested preference priority = PP

Following each request, the sender may maintain its previous rate until a response message is received and, thereafter must shape its forwarding to the new rate received in the response packet.

ARS has no initial rate and so the sender must wait for the first response packet before forwarding data packets.

If no response packet is returned to the sender, as determined following the RESPONSETIMEOUT at the sender, then the request packet must be retried 3 times and then if no response packet is received the FSA signalling should be terminated and the data flow sent without FSA signalling or encapsulation.

7.2.4 ARS clear down

Flow state times out and is terminated if data packets corresponding to that flow are not seen at an FSA node for a period equal or greater than STATETIMEOUT.

An example of message sequences related to FSA transport service set-up and clear down is provided in Appendix II.

7.3 Procedures and format of layer 2 FSA Ethertype packets

7.3.1 FSA Ethertype data packets

The first part of an Ethernet packet to arrive is the 14-byte long Ethernet header. The Ethertype constitutes the last two bytes. Since packets after the Ethertype are normally pictured with four byte widths, the start of the Ethernet header is best pictured starting at byte 3, not byte 1. Data packets have three sections: first the Ethernet header (14 bytes), second the FSA header (4 bytes), and then, the original packet. At the sending FSA signalling edge function, if the Ethertype of arriving packets is not the FSA Ethertype, then a 4-byte FSA header should be inserted between the Ethernet header and the packet body. If the Ethertype of the incoming packet was already the FSA Ethertype, then it will already have an FSA header and no restructuring is required. The FSA header has parameters of Ptype (Packet type), ID1 (first packet ID), and QoS offset (bytes to the end of the packet's identification information). These FSA parameters must then be computed (see clause 7.3.3) and stored into the FSA header. Then the Ethertype in the Ethernet header must be set to the FSA Ethertype. For data packets, the S bit (bit 0 of the FSA header) must be set to zero. The original packet is thereby encapsulated in an FSA header. The format of an FSA data packet is shown in Figure 7-3.

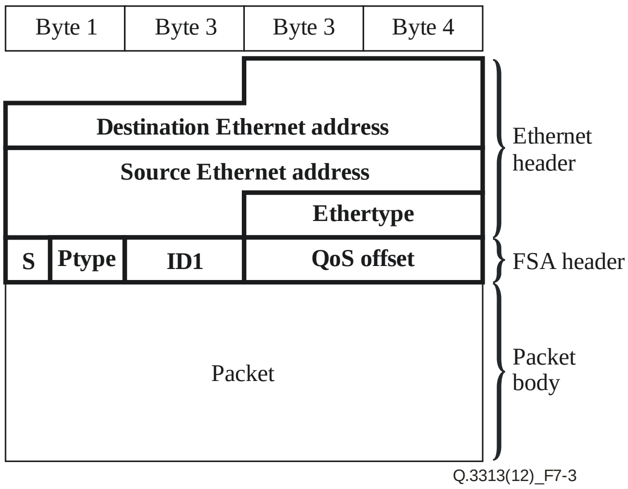

Figure 7-3 – Data packet format

7.3.2 Signalling packets

Signalling packets have four sections: the first is the Ethernet header, the second the FSA header, the third a section of the first data packet of the flow, and the fourth is one or two QoS structures. The terminology, QoS structure, is the name of the 16-byte section which contains the QoS information. Signalling packets are indicated by the S bit, bit 0 of the FSA header, being set to one. Also in a signalling packet the Ptype, ID1 and QoS offset need to be set (see clause 7.3.3). Following the FSA header, signalling packets need to have a section copied from one of the flow's data packets which includes whatever information is needed to identify the flow. The packet section starts at the beginning of the packet and continues for the number of bytes specified by the QoS offset. Right after this portion of the packet follows the first QoS structure. The flow type in the FSA header can be found in the packet. The signalling packets structure is shown in Figure 7-4.

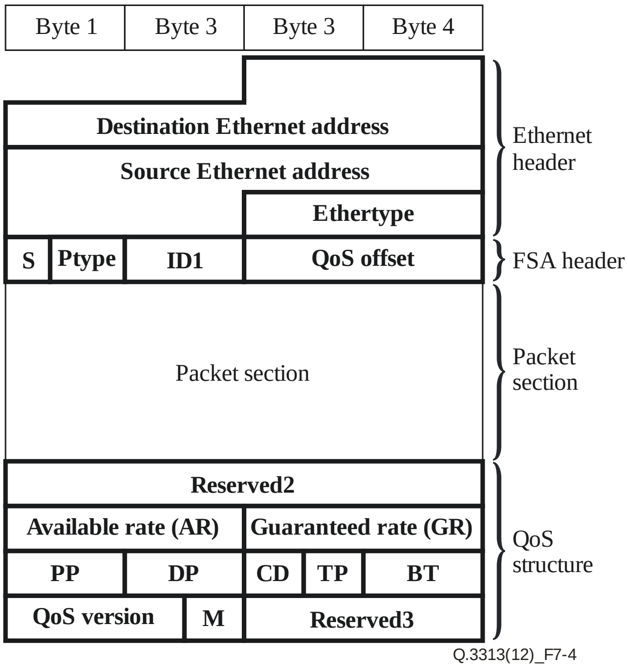

Figure 7-4 – Signalling packet structure

The signalling packet components are:

– **FSA header:**
  - **S**: Signalling bit (4 bits) S = 0x8 for a signalling packet, S = 0x0 for a data packet.
  - **Ptype**: (4 bits) Packet type code as determined in clause 7.3.3.
  - **ID1**: (8 bits) Flow identification, first part as determined in clause 7.3.3.
  - **QoS offset**: (16 bits) Number of bytes between the Ethernet header and the QoS structure.

– **QoS structure:**
  - **Reserved2**: (32 bits) Reserved and set to zero.
  - **Available rate (AR)**: (16 bits) Floating point rate assigned by the network. Floating point format is covered in Annex A.
  - **Guaranteed rate (GR)**: (16 bits) Floating point rate assigned by the network. Floating point format is covered in Annex A.
  - **PP**: (8 bits) Preference priority – Indicates a relative rate priority for ARS and an acceptance priority for MRS flows; 64 levels in the high order 6 bits (bits 0-5), 0 = lowest, 63 = highest. The two low order bits are reserved and should be set to zero.
  - **DP**: (8 bits) Reserved and set to zero (the name was for delay priority, but this parameter has been dropped).
  - **CD**: (4 bits) Change/Direction field – Bit 0: set to zero, Bits 1-3: 0 = No action required, 1 = Request at the start of a flow to negotiate the QoS parameters, 2 = Response returning agreed parameters to the sender, 3 = Reserved, 4 = Reserved, 5 = Renegotiate, a sender request to renegotiate rates on the continuation of a flow, 6 = Reserved, 7 = Close, sent by the sender to close out FSA network nodes state.

– **TP**: (4 bits) Bits 0-1: Reserved and set to zero; Bits 2-3: Type of flow – 0 = Available rate service (ARS, 2 = Maximum rate service (MRS).

– **CH**: (4 bits) Bit 2: Second QoS attached, 0 = single QoS structure, 1 = Second QoS structure follows; Bit 3: Security structure attached, 0 = No security structure follows, 1 = Security structure follows.

– **QoS version**: (12 bits) QoS protocol version field – set to 2.

– **M**: (4 bits) Bit 3: Modified marker. Set to 0 by sender on Request or Renegotiate. Set to 1 by FSA network nodes if any field changed during a request or renegotiate; Bits 0-2: Flow sender depth (FSD) – Set to 1 by sender, increased by one entering a proxy and decremented by one when leaving a proxy.

– **Reserved3**: (16 bits) Reserved and set to 0.

7.3.3 Procedure to determine parameters to store in an FSA header

The sender (the sending FSA signalling edge function) must determine the value of the parameters which will go into the FSA header. The following procedure uses Tables 7-1 and 7-2 to help compute QoS offset, ID1, and Ptype. To simplify the specification of extracting offset data from data packets, three C (content) functions have been defined:

1) C(n) means the contents of byte n. This will be a byte (an 8 bit number).

2) C(n-m) means all the data stored in bytes n through m. This can be from 2 to 32 bytes.

3) The function CQuad(n) means the contents of quad n. A quad is 4 bits so CQuad(n) counts ahead from the Ethernet header by 4*n bits and retrieves 4 bits.

All these functions count from the end of the Ethernet header and include the four byte FSA header which was inserted. The data packet is after the 4-byte FSA header so the functions may retrieve from the FSA header or the data packet. A parameter Q is an offset from the end of the Ethernet header used to "walk" through the FSA header and data packet. Q is then a pointer to the byte now being examined. When the walk through is completed, Q is then saved as the QoS offset which is to be stored in bytes 3 and 4 of the FSA header and will point to the QoS structure if it is a signalling packet. If it is a data packet the QoS offset points to the end of the flow identifier ID2, making it easy for the FSA nodes in the network to find the flow identifier.

The first step is to look up in Table 7-1 the Ethertype that was in the arriving packet and compute Ptype, ID1, and Q based on the formulas in that row.

Table 7-1 – Look-ups based on Ethertype

Ethertype	Ptype	ID1=	Q=
800	1	C(14)	4×CQuad(10)+4
86DD	2	C(11)	44
22EF	CQuad(2)	C(2)	

Once Ptype, ID1, and the initial value of Q are determined, the next step is to finish evaluating ID1 and Q. Normally, this is one step but on occasion the process needs to be repeated several times to determine the correct values of ID1 and Q. This is called a cycle and means that once ID1 and Q are updated the look-up should repeat.

Table 7-2 – Look-ups based on ID1

IF ID1=	New ID1=	New Q=	Next Action
0	C(Q)	Q=Q+8×C(Q+1)+8	cycle
4	C(Q+10)	Q=Q+4×CQuad(2)	cycle
6	ID1	Q=Q+4	done
17	ID1	Q=Q+4	done
41	C(Q+7)	Q=Q+40	cycle
43	C(Q)	Q=Q+8×C(Q+1)+8	cycle
44	C(Q)	Q=Q+8	cycle
50	ID1	Q=Q+4	done
51	ID1	Q=Q+8	done
60	C(Q)	Q=Q+8×C(Q+1)+8	cycle
other	ID1	Q=0	done

Once the look-ups are done, Ptype, ID1, and Q should be saved into the FSA header. Q is saved into QoS offset. To extend ITU-T Q.3313 to other packet types, it is only necessary to augment these tables.

7.3.4 Procedure at the signal origination FSA proxy edge

When a new packet is received at a proxy and the QoS offset computed is zero, it should be encapsulated as a data packet and forwarded. If the QoS offset is not zero, the flow identifiers ID1, ID2 and ID3 (as determined in clause 7.3.5) should be checked against a table of active flows to see if this is a new flow. When a new flow is received at the sending FSA proxy edge, a request type signalling packet needs to be sent before the packet. The signalling packet is created as specified in clause 7.3.2 using a copy of the first packet. A 16-byte QoS structure should then be affixed to the end of the packet section. The service type, rates, and QoS parameters should be set and the CD field is set to 1 for request. The data packet should be held and the signalling packet sent.

After sending the request signalling packet, the sender should wait for the receiver to return a response packet. If a response is not received within the RESPONSETIMEOUT period the sender should resend the request a second and then a third time. If still there is no response, the sender should abort using the FSA protocol and send the data packets of the flow unchanged.

When a response packet is received, the sender should set the sending rate and QoS as returned in the response packet and proceed to encapsulate arriving data packets with the FSA Ethertype and the FSA header as in clause 7.3.1 and send them forward.

7.3.5 Procedure at the FSA network nodes

For each incoming packet which has an FSA Ethertype, an FSA network node should compute the full packet ID if the QoS offset is non-zero to determine the relevant flow. The function C(n) in these computations means the contents of byte n. C(n-m) means all the data stored in bytes n through m. All these count from the end of the Ethernet header and include the four byte FSA header. The full ID consists of three segments; ID1, ID2, and ID3. The FSA header of the packet has Ptype, ID1, and Q=QoS offset. ID2 and ID3 are found as shown in Table 7-3:

Table 7-3 – ID2 and ID3 look-up

If Ptype=	ID2=	ID3=
1	C(17-24)	C(Q-Q+3)
2	C(13-44)	C(Q-Q+3)

The full ID completely defines a specific individual flow. Next, if the FSA header has S = 1 then the QoS structure found at byte Q should be examined. If the CD field of the QoS structure = 1 or 5 (Request or Renegotiate) then the node should adjust the AR or GR rate downward to a rate it can support at the specified preference priority and store the new rate in the QoS structure and mark the M bit = 1 indicating a change was made. The packet should then be forwarded.

7.3.6 Procedure at the signal termination FSA proxy edge

When the receiving FSA proxy edge receives a signalling packet with the FSA Ethertype it first needs to determine if it is nested. Upon receipt of a request signalling packet it should examine the value of FSD in the M field in the QoS structure. If FSD > 1 then the packets of this flow should be passed forward intact unless this is an end-system. A flag should be set to note this for the data packets in the flow. If FSD = 1 then the receiving FSA edge should create a response packet when it receives a request packet. The response packet is the request packet with the Ethernet addresses reversed and the CD field in the QoS structure changed to response (2). The other fields of the QoS structure should remain as received except that the rate can be decreased to the maximum the receiver can support. The same process of returning response packets should be used when receiving renegotiate packets. The signalling packets should not be forwarded. When data packets are received, the Ethertype related to the Ptype should be saved into the Ethernet header, the FSA header removed and the packet forwarded. The Ethertype is found by a look-up into Table 7-4.

Table 7-4 – Ethertype look-up

If Ptype=	Ethertype=
1	800
2	86DD

7.3.7 Renegotiate signalling packets (all services)

After the first 16 packets or 1 second, whichever is sooner, the sender must insert a renegotiate packet to allow the network FSA nodes to adjust the rate. After that, the sender must send a renegotiate packet every one second or every 128 packets, whichever is sooner. The receiving FSA edge should return a response to each renegotiate packet. When a response is received the sender must adjust to the new rate (or less).

8 Security considerations and requirements

8.1 Authentication

Flow state aware user authentication requirements are addressed in [ITU-T Y.2121]. FSA nodes within a domain may authenticate to peer FSA nodes within the domain. FSA nodes communicating as peers across a domain boundary should authenticate with each other. Authentication security requires an additional security data structure which, if bit 3 of the CH field in the QoS structure is set to a one, would follow the QoS structure(s) (see CH in clause 7.3.2). This additional security data structure is for further study.

8.2 Authorization

Flow state aware authorization requirements are addressed in [ITU-T Y.2121].

8.3 Data confidentiality

In the case where user flows with data confidentiality requirements also invoke ARS or MRS, the parameters describing the in-band service request shall not be encrypted.

8.4 Data integrity

Flow state aware parameters may be protected against unauthorized modification while in transit. Flow state aware parameter requests may be protected against replay attacks, in conjunction with data integrity protection binding a set of flow state aware parameters to a specific flow.

8.5 Accountability

It is recommended that flow state aware service invocations are logged, including the identity of the entity requesting the service, the actual service request, and the actual service granted.

8.6 Availability and accessibility

Flow state aware services shall respect the priority preference of each authenticated entity in making admission decisions.

8.7 Privacy

It is recommended that flow state aware services ensure the privacy of user specific policy profiles defining QoS parameter limits and privileges.

8.8 Protection against network attacks, from within or outside

It is recommended that flow state aware services include mechanisms to protect against malformed service invocations and to mitigate denial of service (DoS) attacks.

9 State diagrams for FSA signalling edge functions at origination and destination

FSA signalling edge functions (also called FSA proxies) separate the standard network from the FSA supported bounded subnet. They receive packets of any incoming protocol on the input side, add the FSA signalling packets and send FSA signalled traffic on the FSA side. They also receive FSA signalled traffic from the FSA side, delete the signalling and send the original packets on the other side. Each of these systems contains two processes, a signalling origination process and a signalling termination process. These two processes are shown in clauses 9.1 and 9.2 using a state machine presentation. One difference between an FSA edge function and where the FSA process is installed into the sending computer is that the edge function does not have as much information as to the intent of the user. Thus the FSA edge function should classify available rate flows as ARS and fixed rate flows as MRS. The duration of the flows is not an issue. The rate of MRS flows can often be determined from the information in the incoming packet, but if not, the peak rate may be set to that of a high rate video and since the network does not reserve capacity for MRS flows, there is little harm in specifying a rate higher than needed. The only impact will be that an MRS request may be rejected if it is higher than the remaining capacity in a network trunk. For typical high capacity trunks, this is no problem.

9.1 Origination state machine

FSA signalling is originated by the source of the data flow. The FSA signalling component will maintain a separate FSA signalling origination for each outbound connection. Figure 9-1 shows the four states for the sender.

Signalling origination state machine

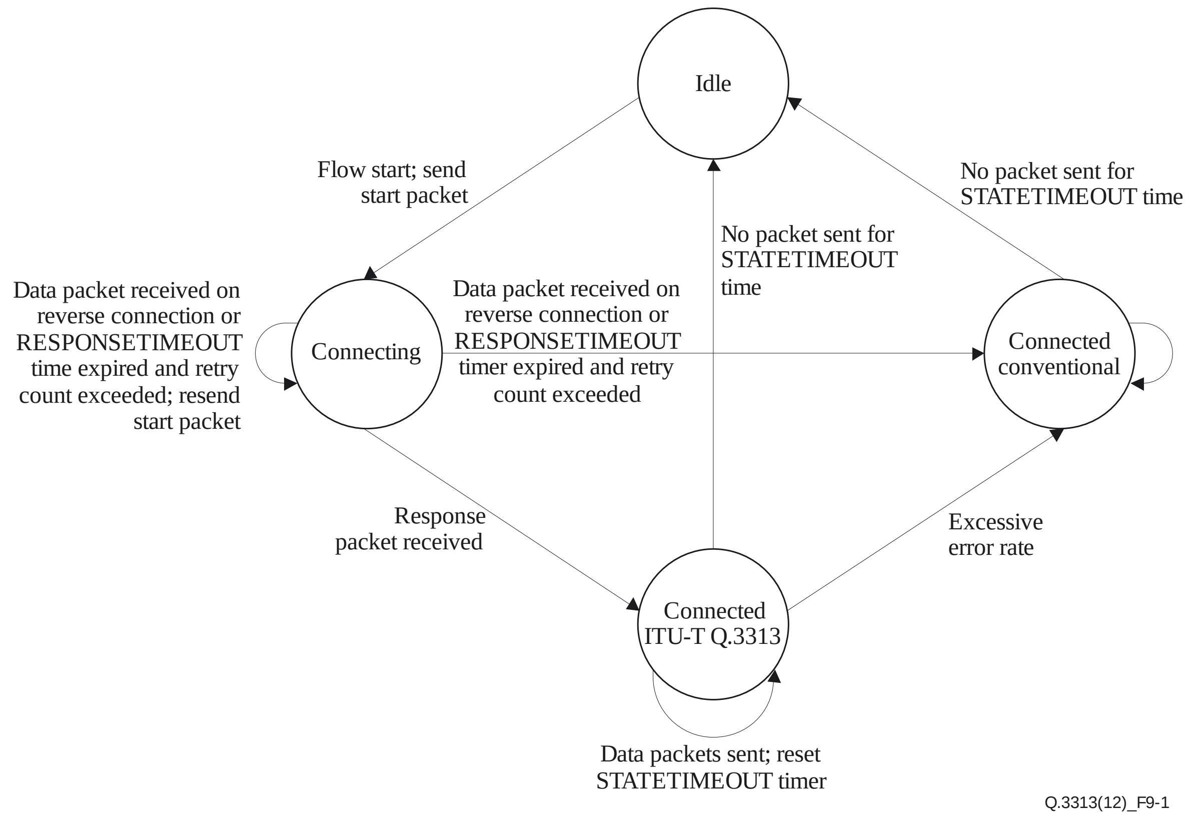

Figure 9-1 – Signal origination state machine

Whenever a packet is received at the input, a full ID1, ID2 and ID1-3 based look-up is used to determine if there is a connection state entry for the flow the packet belongs to. If there is no originating state machine state record then one will be created.

The FSA signalling origination state machine will send a start packet for the new connection and then transition to the connecting state. The FSA signalling origination state machine will remain in the connecting state and resend start packets each time the ITU-T Q.3313 RESPONSETIMEOUT timer expires until the retry limit is exceeded or an ITU-T Q.3313 response packet is received.

9.2 Termination state machine

The transition out of the idle state on the FSA signalling termination state machine is triggered by the reception of a data or start packet for a new flow. When a data packet is received, it is not known whether the packet belongs to an ITU-T Q.3313 flow or a conventional flow. The start packet could have been lost or potentially delayed in the network. The connection state is used to wait for a clear indication that the flow is either an ITU-T Q.3313 flow or a conventional flow. Since all flows are expected to be uni-directional, the origination and termination processes both exist on each system and run independently. Figure 9-2 shows the termination state machine.

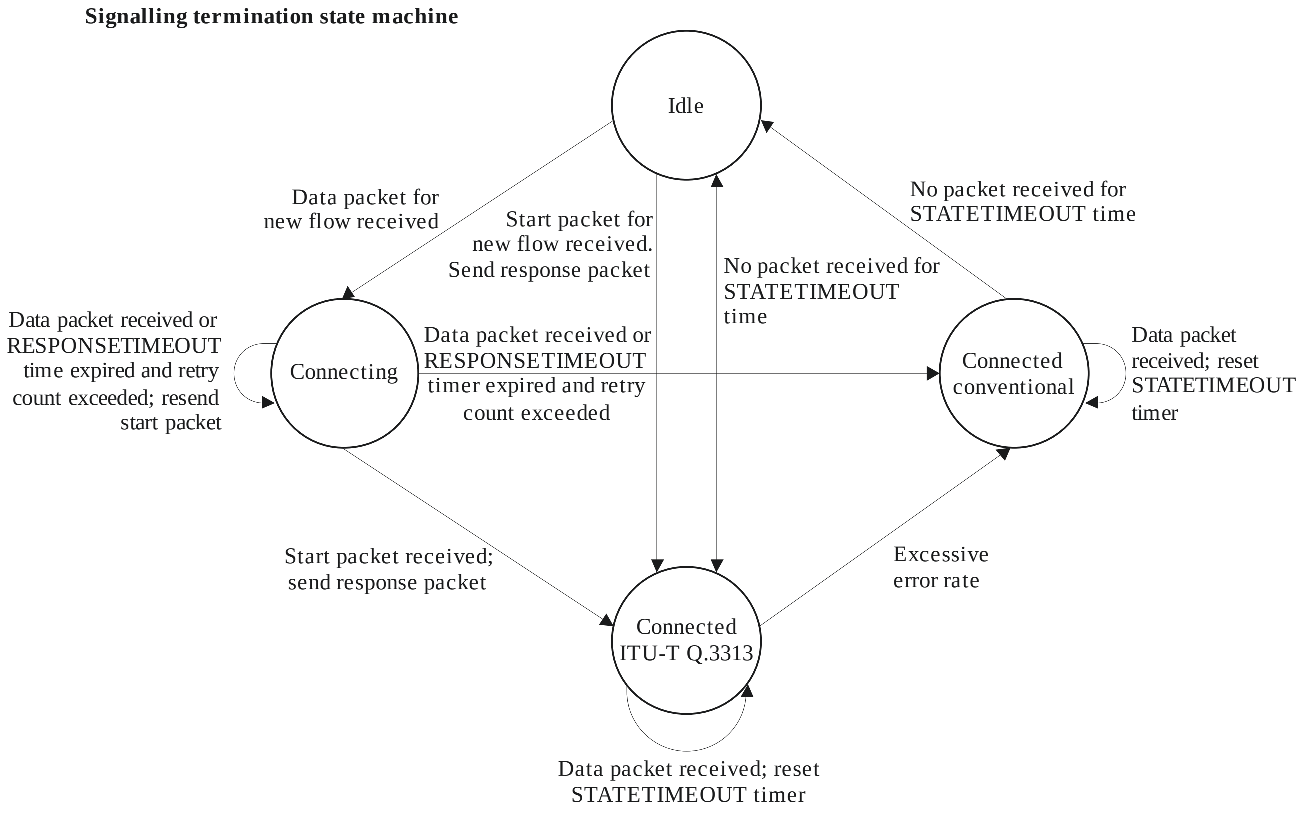

Figure 9-2 – Termination state machine

The connected state for the FSA signalling termination state machine must also deal with reception of additional start packets to deal with the potential loss of response packets. When a start packet is received in the connected ITU-T Q.3313 state, a response packet is sent.

9.3 Signalling renegotiation

To keep the state diagrams readable, renegotiation was not included in either diagram.

Renegotiation is handled in the connected ITU-T Q.3313 states. The FSA signalling origination state machine maintains a renegotiation packet counter and renegotiation timer. The FSA signalling origination state machine sends a renegotiation packet when the renegotiation count is exceeded or the renegotiation timer expires. The default value for the renegotiation packet count is 128 and the default value for the renegotiation timer is one second. The FSA signalling termination state machine sends a response packet every time a renegotiation packet is received.

10 FSA QoS manager

As the signalling packets flow through the bounded subnet, each network node needs to add an FSA flow manager capability which will check the rates and priority requests and if necessary, downward adjusts the requests to the rates which can be supported on the next link. This capability is best shown as a flow chart which acts on each packet passing through. This is shown in Figure 10-1.

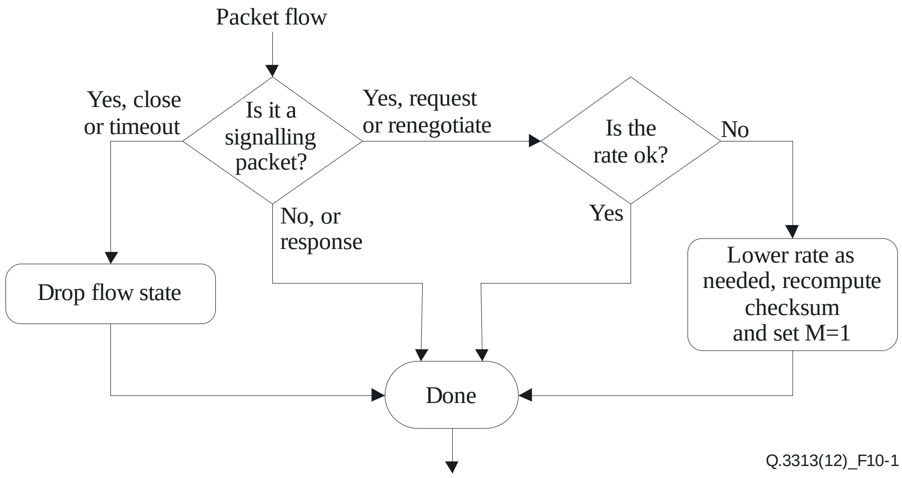

Figure 10-1 – FSA QoS manager flow chart

Annex A

Rules for encoding floating point rates

(This annex forms an integral part of this Recommendation.)

A.1 AR and GR – Encoding floating point

AR and GR fields are encoded according to the following rules:

– The most significant bit of AR and GR is zero and reserved.

– The next bit is nz, where nz = 1 indicates that the number non-zero, and nz = 0 means that the number is zero.

– The next five bits of AR and GR are the exponent E.

– The next nine bits are the mantissa M.

– The rates AR or GR = $(1+M/512)\times2^E$ kilobits per second.

– All zero is interpreted as zero.

– Since E can be as large as 31 and M as large as 511, the maximum rate is 4.291 Tbit/s. The lowest positive rate would be 1 kbit/s.

– This is the same type of floating point number used in asynchronous transfer mode (ATM) except that in ATM its units are cells per second. Since there are no cells in ITU-T Q.3313, the units are kilobits per second.

– The rate is measured using all the bytes in the packet, i.e., the header and the payload. This overhead must be accounted for at each interface so as to compute the trunk load correctly. However, the packet byte count is a usually constant so the user will know what is being asked for and supported. One exception is when fragmentation occurs and in this case those FSA nodes before the fragmentation will have set their rates based on the whole packet, the same as the sender. After fragmentation is done, FSA nodes may see increased traffic due to the fragmentation and thus lower the rate. The lower rate will be enforced. A second case exists if the sender uses header compression. If so, it is the compressed packet size that determines the rate since this is the trunk load incurred across the network. Another case is when header compression is used mid-path to reduce the load on a particular link. Here, the rate marked in the QoS structure should be computed based on the uncompressed header, but the link rate required may be less. Since the average packet size is important to compute the ratio, the first estimate should be based on a best guess assuming large packets, and as the flow proceeds if the packet size is smaller, the extra link capacity can be assigned to other flows.

Appendix I

Mapping FSA services to ITU-T Y.1221 transfer capabilities

(This appendix does not form an integral part of this Recommendation.)

I.1 Maximum rate service (MRS)

The conditionally dedicated bandwidth (CDBW) transfer function specified in clause 8.5 of [ITU-T Y.1221] provides the correct transfer functionality for this service.

Key attributes of this service that determine the underlying transfer capability required are:

Table I.1 – Key attributes of the maximum rate service

ITU-T Q.3313 characteristic	Implication on ITU-T Y.1221 transfer capability
After a response packet, the sender may transmit data packets at any speed up to the requested maximum rate.	Until resources are confirmed as available with a response packet for the flow no QoS commitments can be met.
Only a maximum rate is requested.	After resources have been confirmed this becomes a single token bucket transfer capability (Rp = Max Rate, Bp = the default service provider value). From this point on, all arriving packets that conform to GBRA(Rp,Bp) are conforming.
If a higher rate is required, it must send a new request. Having sent the request, the sender may send at the new rate immediately.	If all packets up to the new maximum rate are considered conforming, then QoS commitments cannot be met until resources at the new rate have been confirmed.

The initial request would signal the commencement of a flow in the "discard first" state with Rp = max rate.

Default values for Bp and the maximum packet size, M, are pre-defined by the service provider rather than requiring to be signalled.

I.2 Available rate service (ARS)

The statistical bandwidth (SBW) transfer function of clause 8.2 of [ITU-T Y.1221] provides the correct transfer functionality for this service.

Key attributes of this service that determine the underlying transfer capability required are to be found in Table I.2.

Table I.2 – Key attributes of the available rate service

ITU-T Q.3313 characteristic	Implication on ITU-T Y.1221 transfer capability
The sender forwards to the receiver a sequence of request packets. All such request packets after the first are labelled as renegotiate packets.	Each flow request could be regarded as a stand-alone transfer capability request.
AR has no initial rate and so the sender must wait for the first response packet before forwarding data packets.	Packets are not conforming until a response has been received from the first request packet that resources are available.

Table I.2 – Key attributes of the available rate service

ITU-T Q.3313 characteristic	Implication on ITU-T Y.1221 transfer capability
QoS parameters – Requested available rate – Highest available rate (per preference priority).	A single bucket GBRA(Rp,Bp), with Rp = Highest available rate and Bp set to the default service provider value. From this point on, all arriving packets that conform to GBRA(Rp,Bp) are conforming. Clause 8.2 of [ITU-T Y.1221] (SBW) supports QoS commitment for conforming packets.
It is a requirement the sender forwards a new AR request after every 128 data packets or 1 sec since it sent the last AR request packet. Following each request, the sender may maintain its previous rate until a response message is received and, must thereafter shape its forwarding to the new rate in the response packet.	Each flow request can be regarded as a stand-alone transfer capability request from one response to the next. From this point on, all arriving packets that conform to GBRA(Rp,Bp) are conforming.

Default values for the highest available rate per preference priority (PP), Bp and the maximum packet size, M, are pre-defined by the service provider rather than requiring to be signalled.

Appendix II

Example of a bounded subnet isolated by proxies

(This appendix does not form an integral part of this Recommendation.)

As the procedures of ITU-T Q.3313 are designed to operate in a bounded subnet, Figure II.1 shows an example of such a bounded subnet where the cable Internet service provider (ISP) has added an end system FSA function (proxy) to its cable modems in the homes connected to one of their points of presence (POPs). Inside the POP before the connection to the Internet core, the ISP installed an FSA node to mark the signalling packet with the best rate for each flow, and terminated the bounded subnet with a high capacity end system FSA function (proxy). Each POP is then a bounded subnet managing the rates of all flows to and from the connected homes and the Internet as they pass through the CMTS equipment, the cable access channels and the link into each home. All FSA signalling is terminated by the home cable modem on one end and the proxy in the POP at the other end.

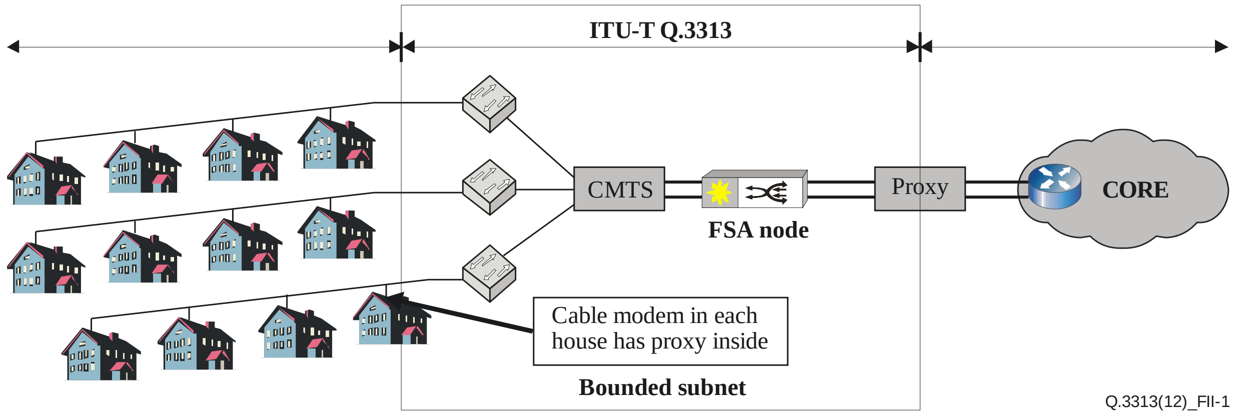

Figure II.1 – Cable ISP operating as a bounded subnet with FSA signalling

In Figure II.1 the traffic in the home is processed in the home modem proxy to encapsulate the traffic and add FSA signalling. The FSA node in the POP marks each flow's signalling requests with a rate that will ensure that no congestion occurs in the cable modem termination system (CMTS), the cable channel, or the home modem. The proxy in the POP de-encapsulates the traffic and eliminates the signalling packets as it passes the traffic to and from the core. This process is totally self-contained as a bounded subnet and its operation will eliminate delay jitter and packet loss across the typically congested cable access path. The CMTS, which typically would control the traffic by discarding many packets causing congestive problems is pre-empted from dropping packets since the FSA control has ensured that no cable or home would be loaded to overload.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Terminals and subjective and objective assessment methods
Series Q	**Switching and signalling**
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2012